

Brad

Trip Report to Kendall Research**Danny Hillis**

On Friday, April 1, 1988, I went with Gordon Bell to visit Kendall Research Corporation, where we met with Henry Burkhardt and one of the other founders (name?). The meeting was open and friendly. At the beginning of the meeting I asked them specifically to not give me any proprietary information, and I did not sign any non-disclosure agreement.

Kendall now is about 50 people, mostly engineers, based at 3 Kendall Square, in the same building as Javelin. They intend to ship a shared memory multi-processor in summer of 1989, performance ranging from 100 to 100,000 MIPS at about \$1,000 per MIP. Their primary intended market is scientific calculation, although they would also like to use the machine for transaction processing. In terms of existing machines, the machine seems to be most directly competitive with Encore and possibly with us at the high end. My own judgement is that the machine will work extremely well in the 100-1,000 MIP range, but not nearly so well in the larger sizes. The company seems adequately funded and has just completed a \$14.1 million financing.

From a software standpoint, the product will be a multi-tasking Berkeley UNIX system, running C and Fortran compilers. The compilers automatically extract parallelism, primarily by unwinding loop and allocate subparts to multiple processors. The operating system will be based on Berkeley UNIX, not on Mach. The one novel software idea seems

to be that the number of processors that a program runs on is determined dynamically at run-time, rather than statically at compile-time, as in the Alliant. Thus, for example, an operating system call may ask for 150 iterations of a loop to be performed and that will be parallelized according to how many processors are available at run-time.

The basic hardware unit is an 8 processor system, with 128 megabytes of memory, dissipating about 400 watts of power. They expect this system to run at 120 MIPS and about 50 linpack megaflops. The price will be about \$150,000. The basic processor is one board, approximately 9 inches x 12 inches, with 20 custom VLSI chips and 512 kilobytes of static RAM for cache. It is a 64 bit processor, with a 64 bit vector floating point unit, and a 32 bit instruction word. It has about 40 integer registers and 64 floating point registers, so that that the total tasks which time is about 150 cycles. They expect to achieve an average throughput of about 15 MIPS. The address space of the machine is 64 bits with segmented protection, and the processor is designed to access memory through a hierarchical caching scheme similar to Encore. The cache is four-way set associative (for data) and two-way for instruction, with process tags. I got the impression that it uses a directory-based cache write-through, although that was the one question that I asked that they wouldn't answer because the data was proprietary. The machine will be entirely air-cooled.

The chips themselves are tab-mounted onto 300 pin leadless chip carriers with pins on 025 centers and reflow soldered onto the surface of the board. (On the tour of the plant I noticed that they had already installed some fairly impressive reflow soldering and wafer probing equipment in a

clean room.) The chips themselves are full custom on a 1.2 micron CMOS process (4.4 metal pitch), with two layers of metal and a low resistance polysilicon. The custom chips are approximately 13 millimeters on a side and there are five custom chip types in the system. The chips are being fabricated and packaged by Sharp, and they expect to initially get about 2 to 3% yield. They have made test chips of basic cells. I noticed that the unmounted packages were stored clean in nitrogen atmosphere chambers, presumably to prevent corrosion on the leads.

They showed me a paper mock-up of the packaging. The boards are packaged with eight boards horizontally in a 19 inch rack on two 4-board backplanes mounted back to back. I didn't see the connectors, but they look like fairly standard BI-type edge connectors. Four I/O boards can also fit in the box. The box dissipates about 400 watts and is air-cooled with fans blowing from front to back. It is about 8 inches high and designed to stack into towers. The box is designed to look exotic, but it would pass our review.

The signal levels on the chips are 5.1 volts internally and 3.1 volts externally. (They say Sharp has experience with this.) Power is distributed to the boxes at 300 volts DC, and goes through a DC-to-DC converter (made by Vicor) and regulated on-board. They felt that this on-board regulation was especially important for the low signal voltages. The 300 volt does mean they need fancier interlocking of the package to pass U.L.

The boards have two groups of signals coming off of them. One is the backplane connections to the shared memory in the group of eight. Another is 128 megabytes per second worth of I/O organized into 16 buffered channels per board. Each 8-processor box has a 1.2 gigabyte per second, 256 bit wide buss, which goes to further caches up the hierarchy.

The chip design was all full custom, based on tools by SCL (?), Mentor, and some of their own. They have fabbed several test chips without problem, but so far none that will be used in the actual machine. Sharp apparently gives them about 5 week turnaround (!). They use SPICE to simulate all of their cells and all clock runs, both on-chip and off-chip. They also run logic simulations in background mode on their 50 Sun workstations.

They are on a very fast schedule and expect working chips by the end of summer, working hardware by fall, beta-test by spring, and shipping by the end of summer. This allows them two turns per chip, max.

My general impression of the place was an intense, high-focused atmosphere, where everyone knew exactly what they were doing. The physical environment was two floors of no-nonsense open offices with a workstation at every desk, and lots of conference rooms with names like "war room." I believe that their schedule doesn't leave much room for mistakes, but my impression is that they will probably make it. I expect the product as conceived will work very well at least in the 100 MIPS to 1,000 MIPS range, although they claim their "design center," size of machine for which they optimized was a 128 processor system, which would be about 2,000 MIPS. Their biggest competitive advantage is their ability to run existing code.